

DEISA – on the way towards an integrated European HPC ecosystem

★ Distributed
★ European
★ Infrastructure for
★ Supercomputing
★ Applications



Participating Sites



BSC	Barcelona Supercomputing Centre	Spain
CINECA	Consortio Interuniversitario per il Calcolo Automatico	Italy
CSC	Finnish Information Technology Centre for Science	Finland
EPCC/HPCx	University of Edinburgh and CCLRC	UK
ECMWF	European Centre for Medium-Range Weather Forecast	UK
FZJ	Research Centre Juelich	Germany
HLRS	High Performance Computing Centre Stuttgart	Germany
IDRIS	Institut du Développement et des Ressources en Informatique Scientifique	France
LRZ	Leibniz Rechenzentrum Munich	Germany
RZG	Rechenzentrum Garching of the Max Planck Society	Germany
SARA	Dutch National High Performance Computing and Networking centre	The Netherlands

DEISA - Essentials



- DEISA has established a systematic network of cooperation of major supercomputing centres in Europe and operates a powerful supercomputing grid on top of national services.
- DEISA provides a collaborative environment for capability computing and data management, and also support for application enabling for complex and extreme computing.
- DEISA paves the way towards usage of the most adequate supercomputer architectures and towards usage of the most powerful supercomputers, for the benefit of European researchers and the advancement of science.
- Many of the services deployed will enable the efficient integration of future European petascale systems into a seamless European HPC ecosystem.

DEISA - Essentials



Contrary to EGEE (the other big European grid project in the EU FP6) the DEISA project is focused on **supercomputing and the specific needs of European computational scientists** in the area of supercomputing.

EGEE: mainly addresses needs for **capacity computing** with standard Linux clusters

DEISA: addresses **capability computing** through the most powerful supercomputers available in Europe.

DEISA is therefore the European counterpart of the American **TeraGrid** project of the US National Science Foundation.

DEISA Key Strategies



- Deployment of global file systems across Europe.
- Unified and seamless access to supercomputing resources throughout Europe
- Europe-wide support for enabling of grand challenge applications
- DEISA Extreme Computing Initiative DECI

Global File System Interoperability Demo between DEISA and TERAGRID during Supercomputing Conference 2005 in Seattle



American and European supercomputing infrastructures linked:
bridging communities with scalable, wide-area global file systems



Involved TeraGrid Sites



Involved DEISA Sites

The DEISA Supercomputing Environment



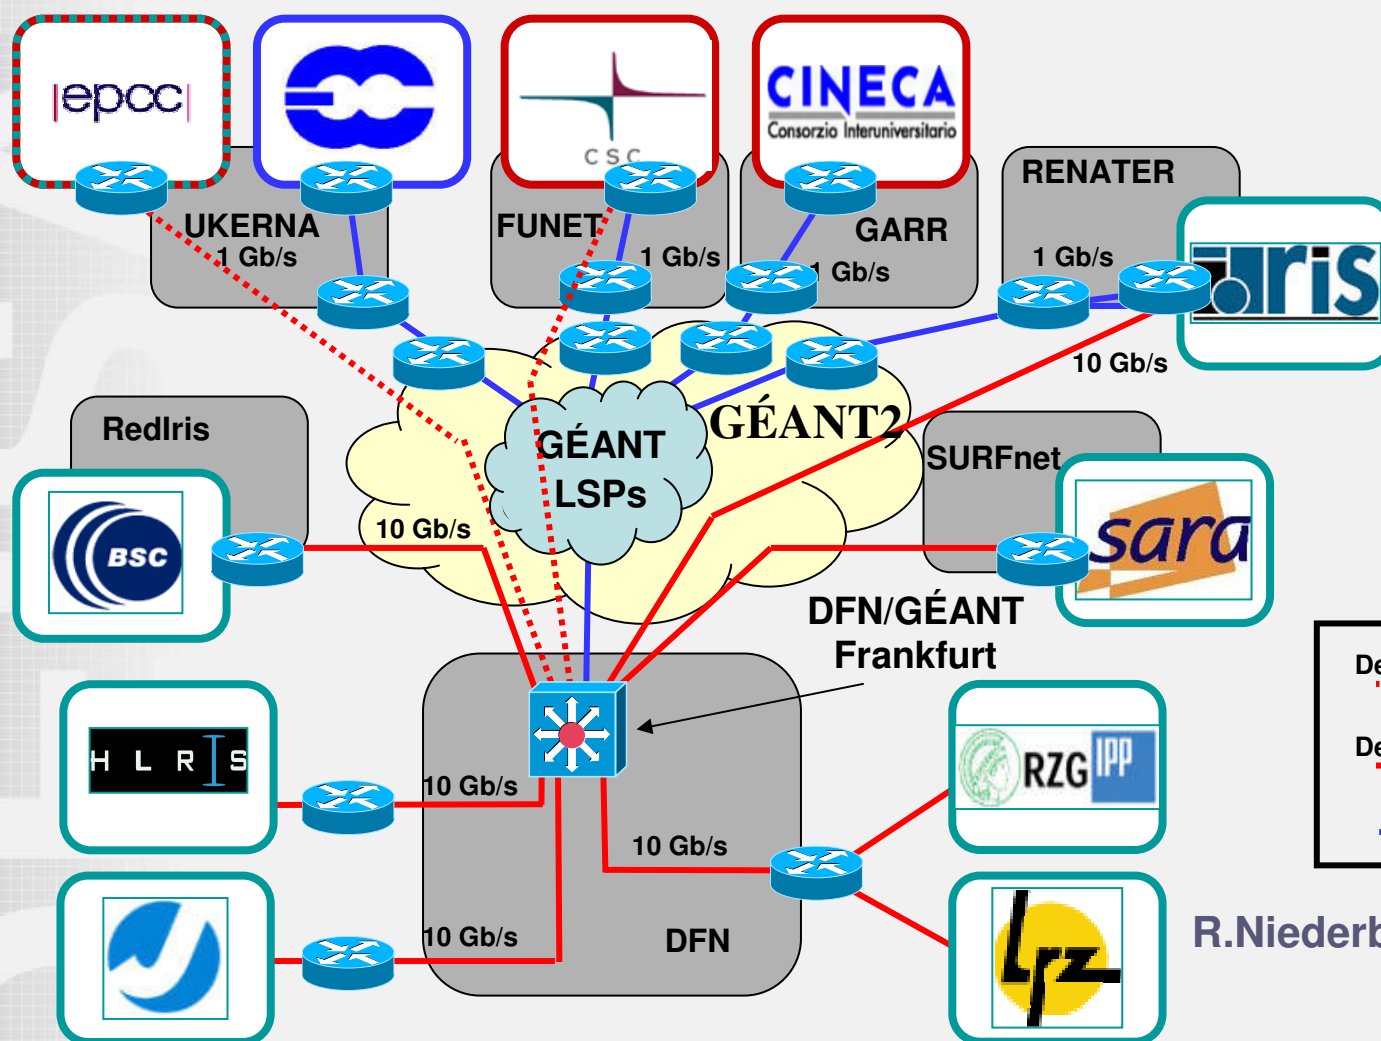
> 30 k processor-cores > 200 TF
of different architectures

- IBM Power4/Power5, AIX
- SGI ALTIX (Itanium2), Linux
- IBM PowerPC, Linux (MareNostrum)
- NEC SX8 vector system, Super UX
- Cray XT4, Linux

***Most Systems interconnected with dedicated 10Gb/s network
– remaining systems upgrading from 1 to 10 Gb/s –
provided by GEANT2 and NRENs***

DEISA - Network

June 2007



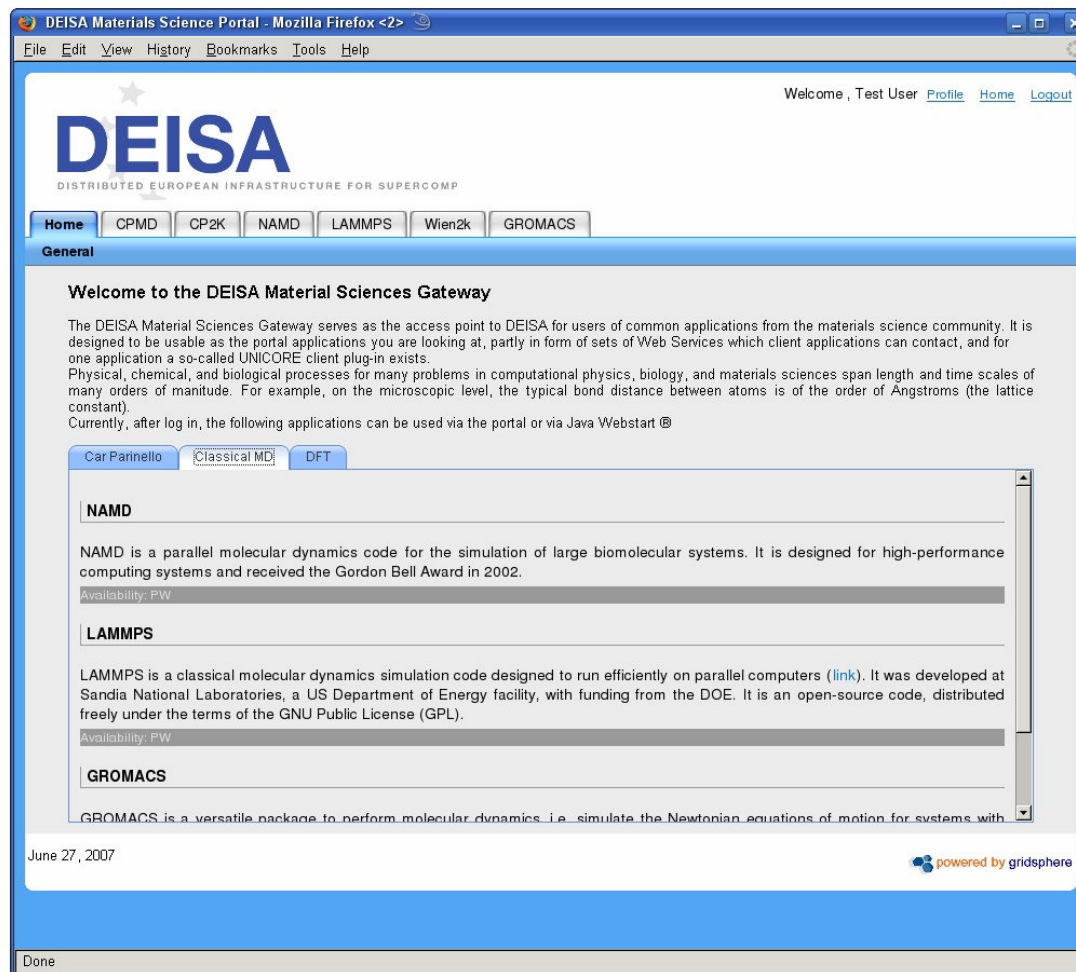
DEISA Service Provision



- Trans-national extension of national supercomputing centres' operational model
 - Service activities based on trans-national working teams with distributed leadership:
 - Network deployment and operation
 - Infrastructure operation
 - Global File Systems
 - Middleware
 - User Support
 - Applications Enabling
 - Security
 - JRA on Middleware R&D
- Joint Research Activities** in
Materials Science, Cosmology, Plasma Physics, Life Sciences,
CFD – industrial applications, Coupled applications
- **Training sessions** (2006 – 2008) in Finland, France, Germany, Italy and Spain

JRA1 in Materials Science

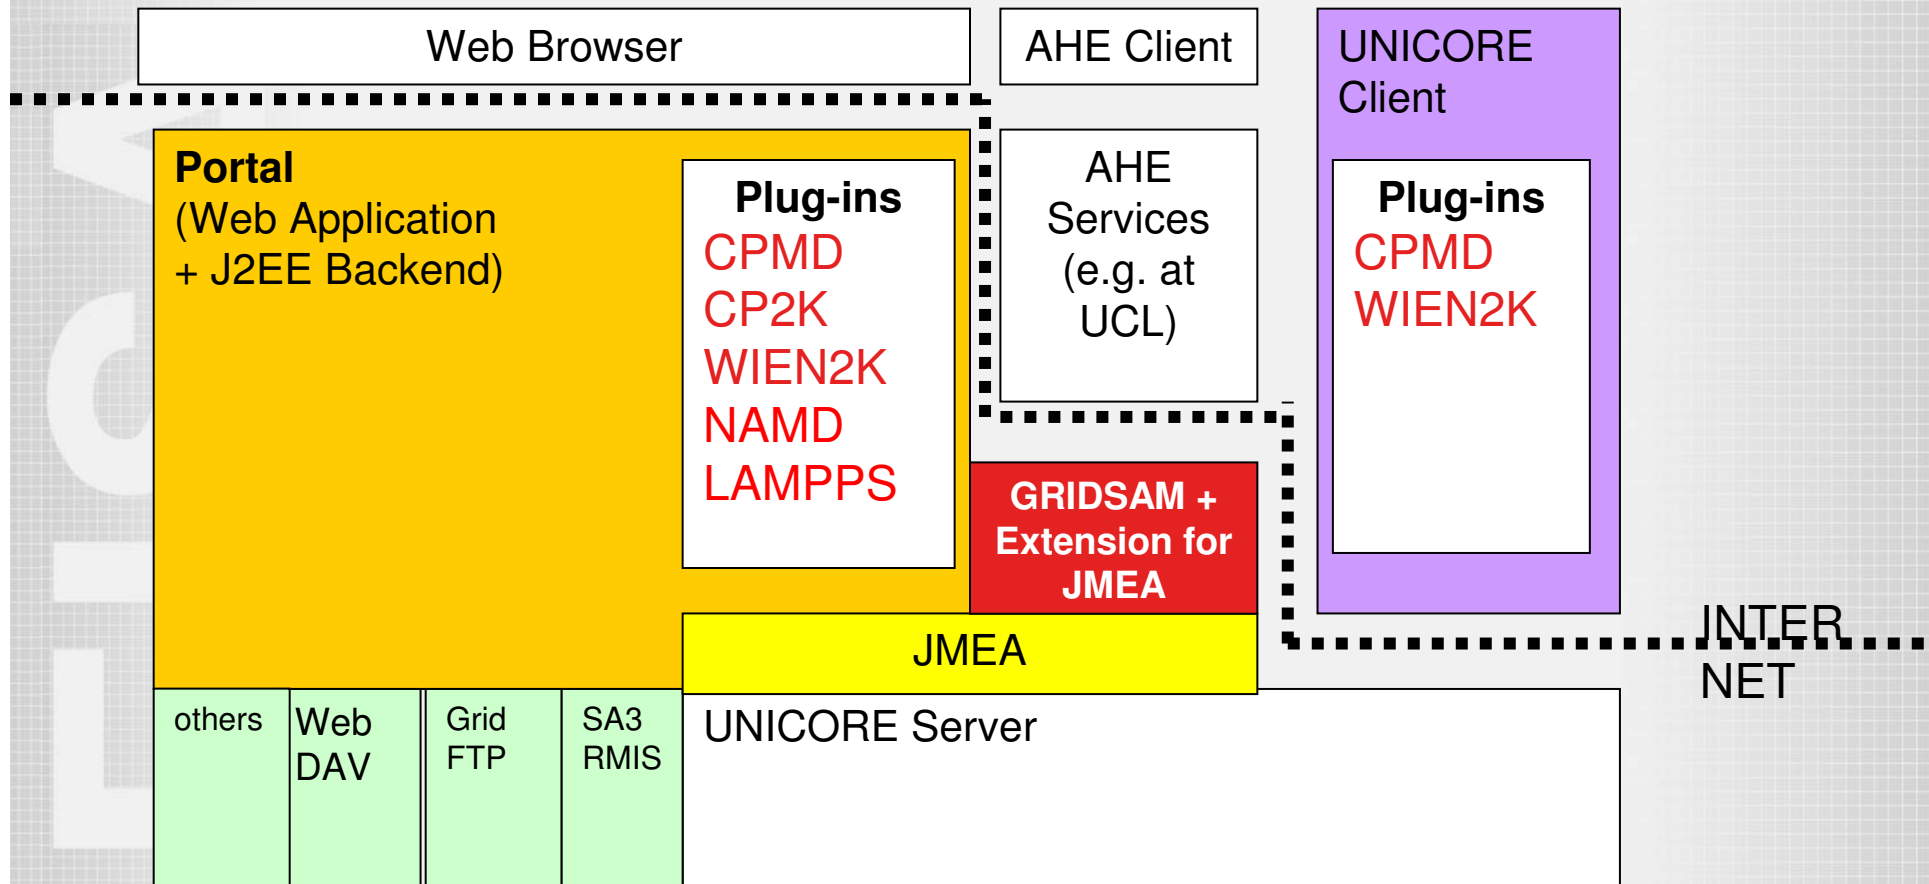
Provision of a set of important materials science simulations codes, grid enabled for DEISA, and integrated in a materials science portal for ease of use



DEISA Materials Science Portal


DEISA Materials Science Portal

- Technical Layout -



- T. Soddemann (2006): *Science Gateways to DEISA*, in: Concurrency Computat.: Pract. Exper., DOI 10.1002/cpe
- T. Soddemann (2007): *Job Management Enterprise Application*, in: W. Lehner et al (eds): Euro-Par 2006 Workshops, Lecture Notes in Computer Science 4375, pp. 254-263, 2007

AHE Demos at SC06 in Tampa, US



GIN at SC'06

Peter Coveney [P.V.Coveney at ucl.ac.uk](http://www.pvcoveney.com)

Sun Nov 12 18:07:38 CST 2006

Over the course of SC'06 we will be demoing the use of the Application Hosting Environment (AHE; <http://www.realitygrid.org/AHE/> and <http://www.omii.ac.uk>) as a single interface to launch and monitor scientific applications (such as CHARMM, NAMD, LB3D, GROMACS, and LAMMPS) on a wide variety of heterogeneous resources provided by the **TeraGrid** and **UK NGS** (running Globus), and **DEISA** (running Unicore).

The demos will mainly consist of launching production scientific simulations investigating the molecular dynamics of HIV-1 protease using the NAMD molecular dynamics code.

Production runs managed by AHE will be running on all three federated grids between 13-17 November 2006. Demos will be given at the following times:

- **Pittsburgh Supercomputer Centre Stand Tuesday 14th - 11:00-12:00**
- **UK e-Science Stand Tuesday 14th - 12:45-13:30**
- **Louisiana State University stand Tuesday 14th 15:00-16:00**
- **EPCC/HPCx Stand Wednesday 15h - 15:30-16:30**

• F. Daniel & T. Soddemann (2007): *Interoperable job submission and management with GridSAM, JMEA, and UNICORE*, Proceedings of GES 2007, (<http://edoc.mpg.de/display.epl?mode=doc&id=316640&col=100&grp=1414>)

JRA2 in Cosmology (EPCC)

Application support for the VIRGO consortium and their application portfolio (including post processing) in the heterogeneous DEISA environment

JRA3 in Plasma Physics (RZG)

Application enabling and support for important codes of the European plasma physics community

JRA4 in Life Sciences (BSC + IDRIS)

- Promote parallel software in the Life Sciences community
- Run big simulations on the DEISA infrastructure that couldn't be done locally
- Provide ease of access to resources
- Application support for life science portal

DEISA Life Science Portal

Distributed
European
Infrastructure for
Supercomputing
Applications

DEISA Life Sciences Portal - SeaMonkey

File Edit View Go Bookmarks Tools Window Help

Back Forward Reload Stop https://biportal.deisa.eu/enginframe/deisa/deisa.xml?_service=BLAST Search Print

Home Bookmarks

DEISA
DISTRIBUTED EUROPEAN INFRASTRUCTURE FOR SUPERCOMPUTING APPLICATIONS

NICE
enginframe

Account: idrtest02 - Project: prtest01

Home Your Jobs Spoolers Logout

The DEISA Life Sciences Portal

Deisa Life Sciences Applications

- BLAST
- NAMD@IDRIS
- NAMD@BSC
- RAXML

BLAST

This service lets users to submit BLAST. Please choose the input parameters:

Blast job name (optional)

Blast program

Blast database

Blast sequence(FASTA,NCBI Accession numbers or GIs)

Expect Value

Filter query sequence (DUST with blastn, SEG with others) ☐ F ☒ T

Query strands to search against database

Produce HTML output ☐ yes ☒ no

Restrict search of database to list of GIs

Use lower case filtering of FASTA sequence ☐ yes ☒ no

Number of concatenated queries

Type of job

DEISA Home

JRA5: Industrial CFD and CAA applications (CINECA)

Exploit DEISA infrastructure capabilities for the use of industry relevant commercial codes, in order to:

- Demonstrate the use of commercial codes on the DEISA infrastructure
- Raise the limit of industrial simulations capabilities a step forward
- Give hints on how to set up commercial codes ASP service into the DEISA infrastructure
- Give a comprehensive overview on use of distributed grid architecture in the field of CFD and CAA

Focus on CFD and CAA applications

JRA6 Coupled Applications (IDRIS)

Example:

3D Combustion / Radiation (FOCUS)

- Study the impact of radiative heat transfer (RHT) on the combustion process (3D – more realistic) in DECI context.
- The 2 coupled codes:
 - AVBP solves Navier-Stokes equations and computes the chemical species evolution.
 - DOMASIUM code computes the radiative field coming from the main species.
- **Achievements**
 - Built and optimized 2 large coupled configurations (load balancing)
 - Gas turbine: emphasize to take account of wall properties (sustain: 350 procs. — 288 + 60)
 - 3D diedra: bring out deep changes in the flame behaviour (sustain: 312 procs. — 192 + 120)
 - Changed the coupling technology with PALM without difficulties
- PhD thesis, 2 publications in international reviews.

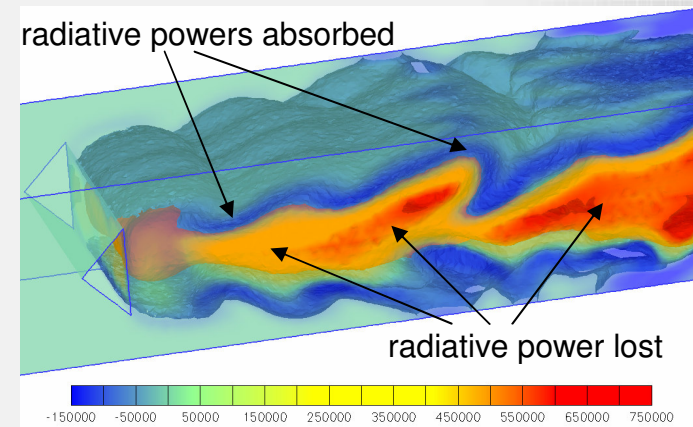


Fig 1: Radiative power in KW.m⁻³




Fig 2: Temperature field in the central plane of the burner ($z = 0$) at $t = 0.55s$

DEISA Extreme Computing Initiative DECI



Performance of challenging supercomputing projects
from all areas of science on DEISA resources

DECI call 2005

- > 50 proposals, 12 European countries involved, co-investigators from US
- > 30 mio cpu-h requested

DECI call 2006

- > 40 proposals, 12 European countries involved,
co-investigators from N + S America, Asia (US, CA, AR, ISRAEL)
- ~ 30 mio cpu-h requested

DECI call 2007

- > 60 proposals, 13 European countries involved, co-investigators from
N + S America, Asia, Australia (US, CA, BR, AR, ISRAEL, AUS)
- > 70 mio cpu-h requested

DEISA - Applications

Consultancy for design or migration of grand challenge applications to the DEISA Grid
optimization and fine tuning of the applications

- Enabling hyper-scaling, data intensive applications, and code optimizations
- Enabling distributed computing, workflows and coupled applications

Establishment of a DEISA benchmark suite with representative applications for different scientific areas, testing essential features of supercomputer architectures

Application enabling / examples (1)



Project SIMU-LU / BSC:

Hyperscaling of GADGET code up to 1500 processors;
I/O improvements for MareNostrum system

Project PSIWAT / CINECA + BSC:

Scalability studies of Quantum-ESPRESSO.
Identification and solution of two performance bottlenecks
improvement of scalability up to at least 512 procs

Project SSSC / CSC + HLRS:

Performance tests of the coupled ECHAM5—MPIOM
in different DEISA architectures;
result: NEC SX8 at HLRS optimal platform;
porting and optimization of the coupled model for NEC SX8

Application enabling / examples (2)



Project TMQCD / FZJ

Compute jobs implemented as workflows within UNICORE for usage on DEISA core sites FZJ, RZG, Cineca, and Idris.

Project DNS-BUMP / IDRIS

Optimisation of MPI Collective Communications;
several topologies and communication schemes were tested
Optimisation of I/O subroutines.

Project SUPERNOVAE / RZG + EPCC

Support for porting the application to the DEISA environment.
Design and preparation of the post-processing which included two steps, one outside of the DEISA infrastructure.
Huge amounts of simulation data had to be made available for post processing for external collaborators located at ITEP, Moscow and at the University Observatory of Stockholm

Application enabling / examples (3)



Project GYROKINETICS / RZG

Hyperscaling of codes to many thousands of processors

Project MolSwitch / RZG + LRZ + BSC + EPCC + CINECA + IDRIS + FZJ



Intensive analysis of CASTEP scalability on various architectures.
Principle limitations detected (and reported to the authors)
for usage of more than 64 to 128 cores per k point.

Turn-around enhancement by:
distribution of MolSwitch project execution over seven DEISA sites.
Five AIX sites CINECA, EPCC, FZJ, IDRIS, RZG addressed via
MC-LL and WAN-GPGS

DECI projects

CHANNEL-2000 (Javier Jiménez):

Flow development under severe adverse pressure gradients.




u velocity near the wall and zoom of the marked zone. Note the high level of detail of the simulation. Visit <http://torroja.dmt.upm.es/~sergio/> for this and other pictures.

- Code tuning for 2048 procs
- Decoupling of computing and I/O (usage of add. 60 procs dedicated to I/O (25 TB))
- Usage of vector instruction set of the PPC970 FX CPU

DECI projects

BET (P. Carloni):

First ab initio investigation of entire electron-transfer proteins in aqueous solution: studies of a copper protein (azurin for *Pseudomonas Aeruginosa*) and two iron-sulfur proteins (rubredoxin from *Clostridium Pasteurianum* and from *Pyrococcus Furiosus*).




Isosurface enclosing the 95% of the total spin density (difference of the up and down electron densities) in the oxidized form of the **rubredoxin from *Clostridium Pasteurianum***.

Optimization of the simulation for big biological systems.

DECI projects

TMQCD (K. Jansen):

Realistic simulations of Lattice QCD, a discretized version of QCD, the theory of the strong interaction.



Dependence of the so-called pseudo-scalar decay constant, f_{PS} , as a function of the quark mass. In the plot, the numerical data are fitted against the prediction of chiral perturbation theory, which is an analytical, effective low energy description of QCD. As can be seen, the fit is excellent and the very precise numerical simulation data allow **determining various low energy constants of chiral perturbation theory with a high, world-record precision.**

Project execution at FZJ, RZG, CINECA, and IDRIS using UNICORE

DECI projects

CAMP (M. Parrinello):

Catalysis by Ab-initio Metadynamics in Parallel




Fig. 1: all-electron and pseudopotential calculations for bulk pyrite





Fig. 2: simulation of the nitrate reduction

DECI projects

SUPERNOVAE (W. Hillebrandt):

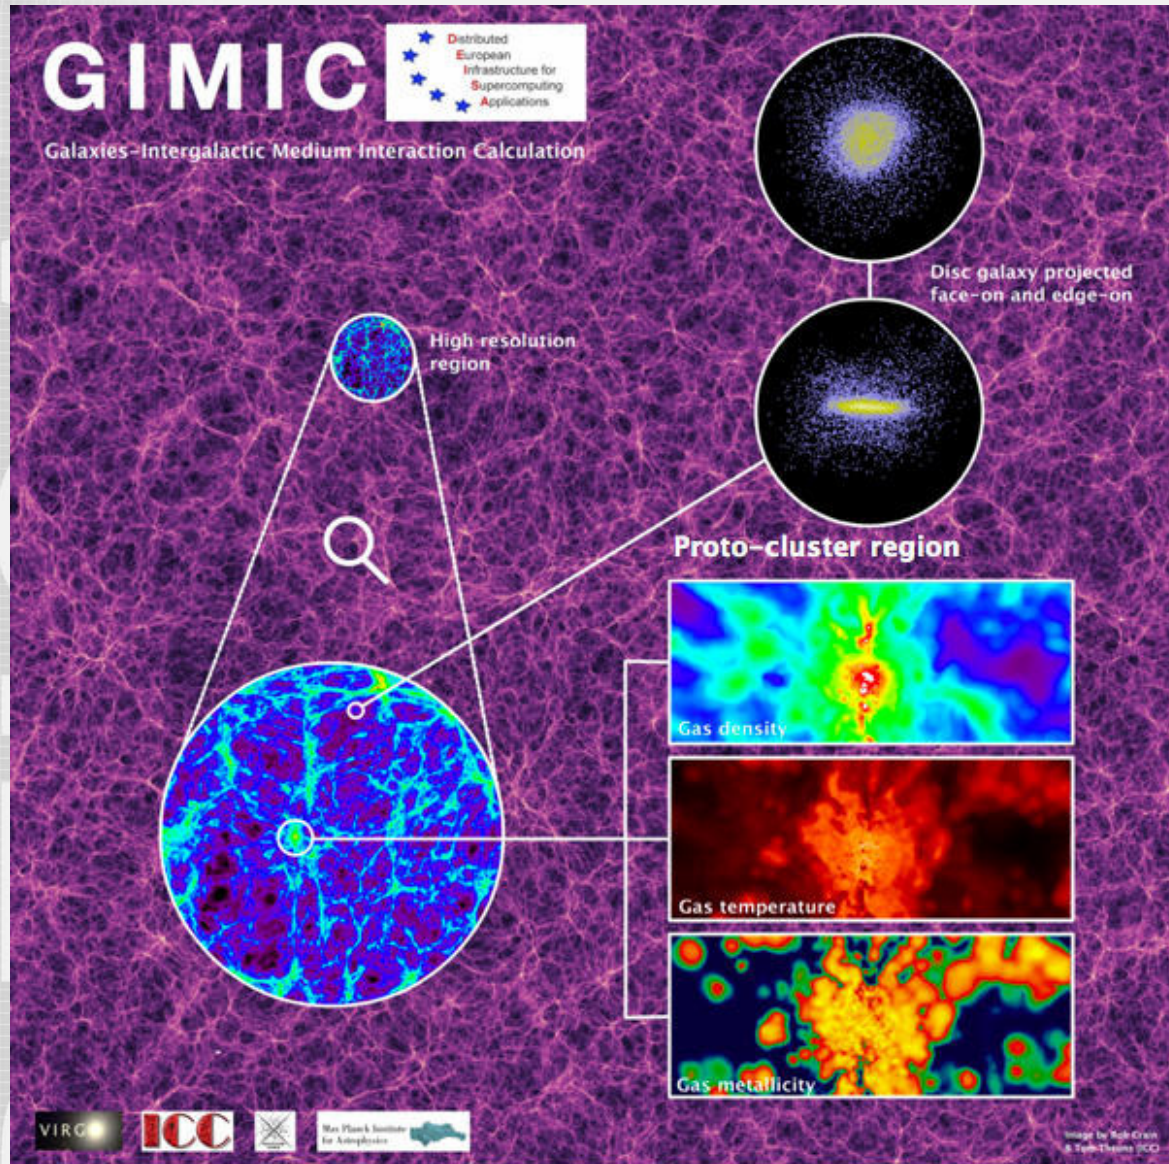
Simulating Type Ia Supernovae on DEISA Supercomputers



Snapshot from a Type Ia supernova explosion simulation carried out within the DEISA framework on the Edinburgh HPCx cluster.

The volume rendering of the logarithm of the density indicates the exploding White dwarf star, and the turbulent thermonuclear flame is represented by the isosurface.

DECI projects



GIMIC

(Simon White and Carlos Frenk):

Galaxies-Intergalactic
Medium Interaction
Calculation

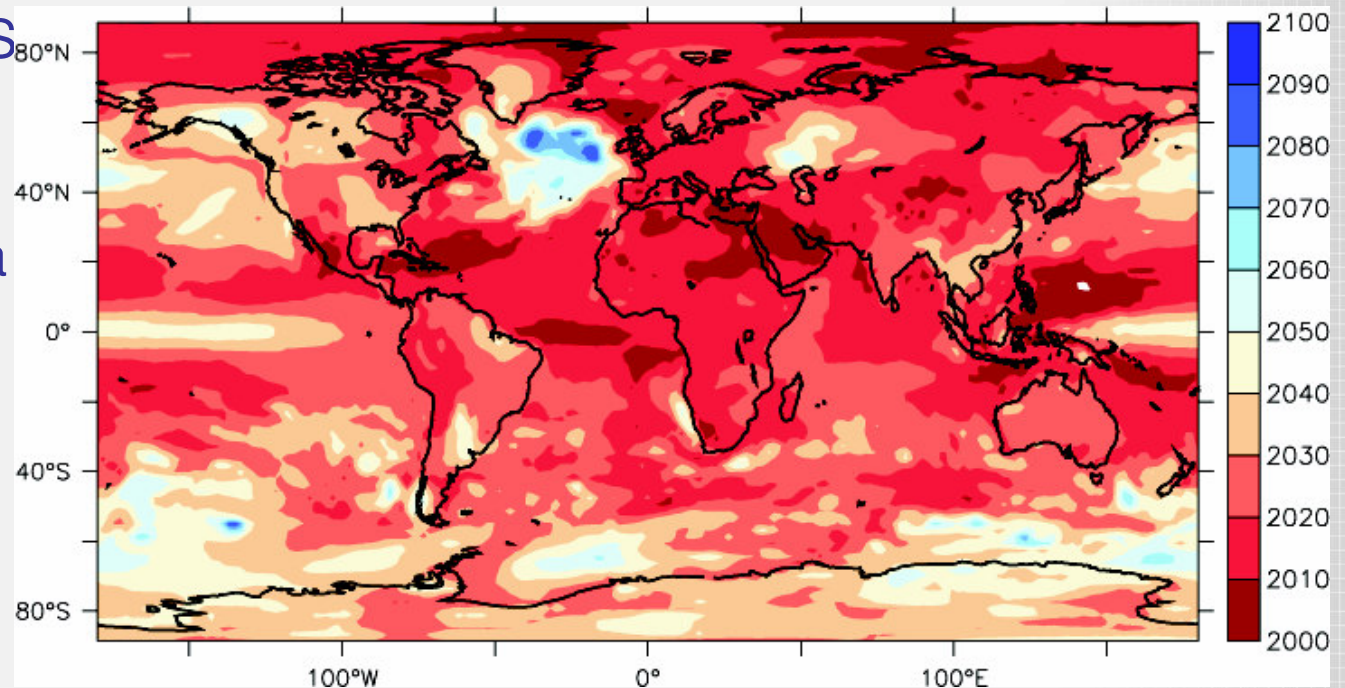
Project of the
Virgo Consortium

DECI projects

Project ESSENCE / SARA + HLRS

Large Climate Simulation by KNMI


- Workflow of simulation runs followed by a postprocessing step
- Coupled simulation using an atmosphere and an ocean model
- Home site SARA,
executed at HLRS
- ~180000 CPU/h
NEC SX-8
- 50 TB output data





DECI projects

POLYRES (K. Kremer)

Curvy membranes make
proteins attractive



- a) proteins (red) adhere on a membrane (blue/yellow) and locally bend it;
- b) this triggers a growing invagination.
- c) cross-section through an almost complete vesicle



DECI-1 Projects - Results



Cover story of Nature May 24, 2007 DECI project POLYRES

Benedict J. Reynwar, Gregoria Illya, Vagelis A. Harmandaris,
Martin M. Müller, Kurt Kremer & Markus Deserno:
***Aggregation and vesiculation of membrane proteins
by curvature mediated interactions,***
NATURE Vol 447|24 May 2007| doi:10.1038/nature05840



**26 publications so far, among them
NATURE, Phys Rev B, Phys Rev Lett, PNAS**

DEISA – Achievements and scientific impact

- Heterogeneous supercomputing infrastructure built-up at continental scale
- Priority high speed end-to-end network
- Common global high performance file system at continental scale to greatly facilitate data management across Europe
- Harmonization of the manifold heterogeneous software environments through the DEISA Common Production Environment (DCPE)
- Uniform infrastructure access through the UNICORE middleware
- Job routing and re-routing across Europe with MC-LL among five AIX platforms to
 - free resources for a big job requiring close to all resources at one site
 - facilitate simultaneous usage of many platforms for independent subtasks of a big project, to accelerate project turn-around in a complementary approach to using UNICORE

DEISA – Achievements and scientific impact / 2

- Portals for transparent access to complex supercomputing environments
- Internal services:
 - system development and deployment of continuous resource and service monitoring,
 - centralized trouble ticket system,
 - accounting facility
 - database for the management of the DECI projects
- High-level application enabling
- Setup of a European Benchmark Suite
- Enabling of cooperative extreme computing in Europe
- Grand Challenge projects performed on a regular basis



The way has been paved!

DEISA – the way has been paved



DEISA



- The way has been paved
- What for?
- Persistence?



PACE

PACE - Partnership for Advanced Computing in Europe -
New Initiative by 14 European countries

Austria, Finland, France, Germany, Greece, Italy, Netherlands,
Norway, Poland, Portugal, Spain, Sweden, Switzerland, UK

Goal: Preparation of the establishment of European
Petaflops centres in the 2010 time frame

The PACE Vision



From cooperative High Performance Computing in Europe

–to

–leadership class

–European HPC-facilities,

–The Large Scale European Supercomputing Infrastructure

integrated in a High Performance Computing Service

– the European Ecosystem -

Partnership for Advanced Computing in Europe - MoU signed April, 17th 2007



DEISA – and the future European HPC ecosystem

PACE planned a 2 year preparational FP7 project
- Proposal submitted in May 2007

DEISA/eDEISA FP6 contracts will end in spring 2008
Continuation planned in FP7 as a three years project DEISA2
- Call open, proposal to be submitted in Sep 2007

Persistence



The general goals of DEISA2 and PACE are in-line with the objective of the ESFRI roadmap

ESFRI = European Strategy Forum on Research Infrastructures

DEISA2 and PACE will jointly contribute to building up the European HPC ecosystem

Complementary contents in DEISA2 and PACE will be coordinated

INFRA-2007-1.2.3: e-Science Grid infrastructures

Expected Outcome

- Providing persistent, cross disciplinary services to pan-European virtual research communities, with increased levels of trust and confidence
- Sharing more computing resources (including distributed supercomputing), scientific instrumentation and data facilities across multiple scientific disciplines
- Promoting open standards and interoperability
- Innovating and consolidating middleware technology
- Adopting robust, reliable and scalable authorisation and authentication schemes
- Implementing attractive mechanisms for the pooling of further resources across a very broad range of user communities
- Evolving towards the adoption of more sustainable organisational models for the provision of services

INFRA-2007-1.2.3: e-Science Grid infrastructures

Expected Impact

- Maintaining the world class performance of European Grid infrastructures and their global relevance
- Creating a new generation of e-Infrastructure services and new work and business models for science and engineering
- Fostering the use of grids by new research communities
- Further integrating national initiatives and increasing international collaboration

Examples of activities

- Establishing a production-quality scalable, reliable infrastructure, building on the current achievements and being accepted on a global scale
- Addressing issues of an open market for grid services (e.g. implementing incentive-based models for providing grid resources, deploying relevant accounting schemes, etc.)

DEISA

-

Paving the way was well worth the effort!